OptEx: A Deadline-Aware Cost Optimization Model for Spark

Subhajit Sidhanta*, Wojciech Golab†, and Supratik Mukhopadhyay*
*Louisiana State University, Baton Rouge, Louisiana, USA, Email: {ssidha1, supratik}@csc.lsu.edu
†University of Waterloo, Waterloo, Ontario, Canada, Email: wgolab@uwaterloo.ca

arXiv:1603.07936v1 [cs.DC] 25 Mar 2016

***Abstract*—We present OptEx, a closed-form model of job execution on Apache Spark, a popular parallel processing engine. To the best of our knowledge, OptEx is the first work that analytically models job completion time on Spark. The model can be used to estimate the completion time of a given Spark job on a cloud, with respect to the size of the input dataset, the number of iterations, the number of nodes comprising the underlying cluster. Experimental results demonstrate that OptEx yields a mean relative error of 6% in estimating the job completion time. Furthermore, the model can be applied for estimating the cost optimal cluster composition for running a given Spark job on a cloud under a completion deadline specified in the *SLO* (i.e., Service Level Objective). We show experimentally that OptEx is able to correctly estimate the cost optimal cluster composition for running a given Spark job under an SLO deadline with an accuracy of 98%.**

I. Introduction

Optimizing the cost of usage of cloud resources for running data-intensive jobs on large-scale parallel processing engines is an important, yet relatively less explored problem. Cloud service providers, like Amazon, Rackspace, Microsoft, etc., allow users to outsource the hosting of applications and services to a cloud using clusters of *virtual machine instances*. The cloud service providers charge a *service usage cost* to the user on the basis of the hourly usage [1] of the virtual machine instances. The cloud service providers present the users with a variety of virtual machine instance types to choose from, such as micro, small, large, etc., for Amazon Ec2 [1]. Each virtual machine instance type has a different specification, in terms of CPU, I/O, etc., and different hourly usage cost. The *cost optimal cluster composition* specifies a number of virtual machine instances (of different virtual machine instance types), that enable execution of the given job under the SLO (i.e., Service Level Objective) deadline, while minimizing the service usage cost. However, the current state-of-the-art [2, 3, 4] cluster provisioning solutions do not ensure that a given SLO deadline for job execution is satisfied, while at the same time above service usage cost is minimized.

We present OptEx[1], a closed-form job execution model for Apache Spark [5], a popular parallel processing engine. OptEx can be used to determine the cost optimal cluster composition, comprising virtual machine instances provided by cloud service providers, like Amazon, RackSpace, Microsoft, etc., for executing a given Spark job under an SLO deadline. As far as we know, OptEx is the first work that analytically models job execution on Spark. OptEx analytically models the *job completion time* of Spark jobs on a cluster of virtual machine instances. It decomposes the execution of a target Spark job into smaller phases, and models the completion time of each phase in terms of: 1) the cluster size, the number of iterations, the input dataset size, and 2) certain model parameters estimated using *job profiles*. OptEx categorizes Spark applications into application categories, and generates separate job profiles for each application category by executing specific *representative jobs*. The model parameters for the target job are estimated from the components of the job profile, corresponding to the application category of the target job. Experimental results demonstrate that OptEx yields a mean relative error of 6% in estimating the job completion time.

Using the model of job completion time (OptEx), we derive the objective function for minimizing the service usage cost for running a given Spark job under an SLO deadline. The cost optimal cluster composition for running the target Spark job under the SLO deadline is obtained using constrained optimization on the above objective function. Experimental results demonstrate that OptEx is able to correctly estimate the cost optimal cluster composition for running a given Spark job under an SLO deadline with an accuracy of 98%. We also demonstrate experimentally that OptEx can be used to design an optimal schedule for running a given job on a given cluster composition under an SLO deadline.

Consider the use case where a web development company needs to run a Spark PageRank application to determine the most important web pages they developed over the years, using the infrastructure (cluster) provided by a popular cloud provider, like Amazon. Using state-of-the-art [2, 3, 4] prior experience-based provisioning techniques, the company may provision a cluster of 30 m2.xlarge Amazon Ec2 instances to run the Spark job, under an SLO deadline of 70 hours. In this case, they may end up actually finishing the job in 40 hours, incurring a service usage cost of $168.45 (at the hourly rate of 0.1403 using the pricing scheme from [1]). However, with OptEx, the job would have completed in 60 hours using only 10 m2.xlarge Amazon Ec2 nodes, incurring just $84.18, while satisfying the deadline. Thus, OptEx helps minimizing the service usage cost without violating the SLO deadline.

The technical contributions of this paper are summarized as follows.

- We present OptEx, an analytical model for Spark [5] job execution.
- We provide a technique for estimating the cost optimal cluster composition for running a given Spark job

[1]The project is partially supported by Army Research Office (ARO) under Grant W911NF1010495. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the ARO or the United States Government.

under an SLO deadline, using the above model.

A. Motivation

Scaling out (i.e., adding nodes to the cluster) [6] is the common way of increasing the performance of parallel processing on the cloud. Cost is an important factor in scaling out, with the cost of cluster usage increasing linearly with the number of virtual machine instances in a cluster, evident from analysis of the Amazon Ec2 pricing policy [1]. Hence for minimizing the service usage cost, determining the optimal cluster size for executing a given job is of utmost importance. In the current state-of-the-art [2, 3, 4], a cloud service consumer can choose the required cluster configuration in one of the following manners: 1) arbitrarily, or 2) make an informed decision using previous experience of running similar jobs on the cloud. However none of the above strategies ensure that the SLO deadline is satisfied, and at the same time the service usage cost is minimized. Elastisizer [6] is the only successful work in this direction, but it addresses Hadoop MapReduce and does not address Spark.

B. The OptEx Approach

OptEx decomposes a Spark job execution into different phases, namely the initialization phase, the preparation phase, the variable sharing phase, and the computation phase. Following an analytical modelling approach, OptEx expresses the execution time of each phase in terms of the cluster size, number of iterations, the input dataset size, and certain model parameters. Similar to the ARIA framework [7], which applies profiling for scheduling Hadoop MapReduce jobs, OptEx estimates the model parameters with the components of the specific job profile corresponding to the application category of the target job. ARIA uses Hadoop-specific parameters for profiling [7], and hence is unsuitable for application to Spark. OptEx considers job completion deadline as an SLO parameter [8, 9], that acts as the constraint for minimizing the service usage cost. An objective function for the service usage cost is obtained based on the job execution model. Constrained optimization techniques [10] are applied on the objective function to estimate the cost optimal cluster composition for finishing a Spark job within a given SLO deadline.

II. Spark Job Execution Phases

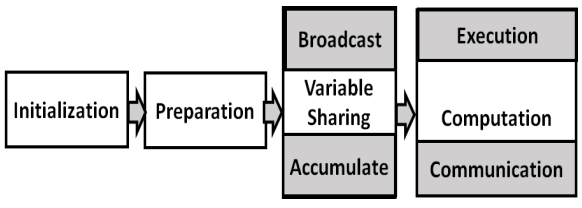

Fig. 1: Phases in a Spark Job Execution Flow

We decompose a typical Spark job execution flow into logically distinct phases illustrated in Figure 1. Each of these phases behave differently with respect to variations in the number of iterations, the cluster size, and the dataset size. The first phase in a Spark job is the *initialization* phase, which performs activities like class loading, symbol table creation, object initialization, function loading, and logger initialization. The second phase is the *preparation phase*, which is responsible for job scheduling, resource allocation, and context creation. The initialization and preparation phases are relatively invariant to changes in input variables [5]. The next phase is the *variable sharing phase* that deals with broadcasting or accumulating blocks of data from the Spark master to the workers.

Spark uses a novel in-memory data structure called the *RDD* (i.e., resilient distributed dataset) for fast and fault tolerant computation [5]. Internally, each Spark job is processed as a permutation of several *unit RDD tasks* (operations like flatmap, reduce, etc), that are executed in parallel on the worker nodes. Spark provides a wide range of built-in unit RDD tasks, packaged within several library modules [11], like MLlib, Spark SQL modules, etc. During the last phase, i.e., the *computation phase* (Figure 1), the given application makes calls to methods from the above library modules, which in turn triggers the respective unit RDD tasks on the workers. The computation phase comprises: 1) the *communication phase* that communicates the intermediate variables among the workers, and 2) the *execution phase* that involves the actual execution of the unit RDD tasks on the workers. The lengths of the variable sharing phase and the computation phase monotonically increase with the input variables, i.e., the number of iterations, the cluster size, and the dataset size [5]. In particular, the variable sharing phase and the computation phase are repeated under iterations, and the lengths of the above phases increase with respect to number of iterations.

III. Application of Profiling For Estimating the Model Parameters

The most common technique for estimating the performance of a given job [7, 12] is using a standard profiling tool, that measures real-time performance statistics, to generate job profile for the target job using a representative job. OptEx categorizes Spark applications, and uses profiling to generate separate job profiles for each application category with representative jobs for each category. Components of the job profile are used as estimates for the model parameters of the target job.

A. Application Categorization

As discussed earlier, OptEx categorizes Spark applications into application categories, chooses a representative job for each category, and creates the job profile for that category using the respective representative job. Application categorization is a difficult open problem, dependent on the application domain, and beyond the scope of this paper. OptEx allows the developers to choose their own categorization scheme.

In this paper, we categorize Spark applications according to the category of library modules that an application uses. Each Spark application uses specific libraries depending on the business logic [5]. The Apache Spark distribution [11] currently organizes the library modules into following four categories: 1) Spark SQL, which supports Apache Hive or JDBC queries, 2) Spark Streaming, which comprises streaming applications, 3) MLlib, which supports machine learning, and 4) GraphX, which facilitates working with graphs and collections. Thus, OptEx uses four application categories, and

a specific job profile for each category is obtained using a representative job for each category.

B. Choice of Representative Jobs For Each Category

The execution phase (Figure 1) of a Spark job comprises a permutation of low-level unit RDD tasks. We call an application a to be a representative job for a given job j, if: 1) the job a contains all the unit RDD tasks comprising the job j, and 2) if job j is iterative, a is also iterative, and vice versa. According to the given categorization scheme (Section III-A), OptEx categorizes applications on the basis of the Spark library modules they use. The Apache Spark distribution web page [11] describes an example application for each group of library modules. For a given application category, the respective example application is chosen as the representative job, under the given categorization scheme. By the design of the Spark libraries, these chosen jobs trivially satisfy the above two conditions for being a representative job.

The custom Spark application, mentioned as an example in the web page of Spark Streaming library [11], is used as the representative job for the applications using the Spark Streaming. It runs on the Twitter dataset [13] , and lists the current tweets on a sliding window. Similarly, the representative application for the MLlib group of applications is the movie rating application MovieLensALS [11]. The input workload for the MovieLensALS applications is the MovieLens dataset made available by Netflix at grouplens.org [14]. PageRank is the representative application for the GraphX group of applications [11]. In the absence of an example application for the Spark SQL category, the widely used Big Data Benchmark [15] developed by AMPLab is used as the representative job for this category. It has been widely accepted as a benchmark for Spark SQL jobs [11].

C. Estimation of Model Parameters from the Job Profile

TABLE I: Glossary of symbols and terms

T_{vs}	Estimated completion time for the variable sharing phase		T_{Est}	Estimated job completion time
n	The cluster size		T_{init}	Estimated completion time for the initialization phase
M_a^k	Execution time of the k^{th} RDD operation for job a		T_{prep}	Estimated completion time for the preparation phase
T_{Rec}	Recorded execution time		$iter$	Number of iterations
t	number of possible instance types		s	size of input dataset
T_{comp}	Estimated completion time for the computation phase		T_{commn}	Estimated completion time for the communication phase in T_{comp}
$T_{vs}^{baseline}$	baseline value of T_{vs}		$coeff$	coefficient of T_{vs} in T_{Est}
$T_{commn}^{baseline}$	baseline value of T_{commn}		cf_{commn}	coefficient of T_{commn} in T_{Est}

A Spark job is typically written in a high-level language (like Scala, Python, etc.) internally executed in different phases (Figure 1). The length of the initialization phase T_{Init} and the length of the preparation phase T_{prep} remain constant to variations in the input variables [5]. The length of the execution phase (T_{exec}) and the length of the variable sharing phase (T_{vs}) increase monotonically with respect to the input variables [5]. Thus, the length of these phases in the execution of a representative job (contained in the job profile) can act as the point of reference, i.e., *baseline*, for measuring the length of the corresponding phases in the target job. In this section, we elaborate how these baseline values in the job profile can be used for estimating the parameters of the model for each job phase.

TABLE II: An Example Job Profile: Profile for MLlib jobs on m1.large instances

App	T_{init}(sec)	T_{prep}(sec)	$T_{vs}^{baseline}$(sec)	$coeff$	$T_{commn}^{baseline}$(sec)	cf_{commn}	T_{exec}	
							RDD task	M_a^k(ms)
ALS	20	13	15	0.004	11	0.07	mean	100
							map	98
							flatmap	72
							first	5
							count	124
							distinct	300

During profiling, the representative application a is run on a single node, and the length of the initialization phase, the preparation phase, the variable sharing phase, and the communication phase (Figure 1) is recorded in the job profile. The length of the above phases in the job profile act as baseline values for estimating the lengths of the corresponding phases in a given target job. The length of the initialization phase T_{Init} (Table II) and the length of the preparation phase T_{prep} for a given job are directly estimated from the lengths of the corresponding phases in the job profile (since, as discussed in Section II, these phases remain constant with respect to the variations in the input variables). As elaborated in Section II, the length of the variable sharing phase T_{vs} increases monotonically [5] with respect to the cluster size and the number of iterations. Hence, T_{vs} is expressed as a function of:

- The input variable n represents the number of nodes.
- The input variable $iter$ represents the number of iterations.
- The baseline value $T_{vs}^{baseline}$, contained in the job profile, representing the length of the variable sharing phase of the representative application. It is the baseline for estimating the length of the variable sharing phase T_{vs} of a given target job.

The increase in the length of the variable sharing phase T_{vs}, estimated relative to the baseline $T_{vs}^{baseline}$, in terms of the given values of the input variables n and $iter$ (we compute the total duration of the variable sharing phase across all iterations), i.e., T_{vs} is expressed as:

$$T_{vs} = coeff \times iter \times n \times T_{vs}^{baseline}, \quad (1)$$

where $iter$ is the number of iterations, n is the number of nodes, $T_{vs}^{baseline}$ is the baseline value, and $coeff$ is a coefficient term. The coefficient term $coeff$ is empirically estimated during job profiling using curve fitting on the results of repetitive experiments with the representative job. The length of the computation phase T_{comp} (Table I) is made up of two logical components: the length of the communication phase T_{commn}, and the length of the execution phase T_{exec} (Figure 1).

The communication phase is responsible for fetching the values of the intermediate variables computed by tasks in the earlier stages of the given job. While profiling, the length of

the communication phase of a representative application a on a single node is recorded in the job profile as $T_{commn}^{baseline}$. It serves as the baseline measure against which the length the communication phase T_{commn} is estimated. The size s of the input dataset is given in bytes (for example, the size of the input for the wordcount application is given as the size of the input files). Since the length of the communication phase T_{commn} increases with respect to the input variable s [5], T_{commn} is expressed as a product of the input dataset size s, a coefficient cf_{commn}, and the baseline value $T_{commn}^{baseline}$, where cf_{commn} and $T_{commn}^{baseline}$ are obtained from the job profile. Again, the coefficient cf_{commn} is empirically estimated in the profiling stage applying curve fitting on the outputs of experiments with the representative job. Thus,

$$T_{commn} = cf_{commn} \times T_{commn}^{baseline} \times s \tag{2}$$

As discussed in Section III-B, the execution phase of a Spark job comprises a permutation of unit RDD tasks. The OptEx job profile records the average running time M_a^k (Table I) of each unit RDD task component k comprising the representative Spark application a. If there are multiple occurrences of an RDD task i in a, we consider the average running time for all occurrences of the task i. By design the representative job for an application category contains all the unit RDD tasks comprising any given job in that category. Hence, the length of the execution phase of the given Spark job is estimated as a function of the average running time M_a^k of each unit RDD task k in the job profile.

IV. Derivation of the Spark Job Execution Model

A. Input Variables

OptEx accepts the following input variables: the size s of the input dataset in bytes, the number of nodes n in the cluster, and the number of iterations $iter$ in the given job [5]. The number of iterations for an iterative Spark application is typically passed as a runtime argument by the developer [5]. Moreover, Spark applications typically have only few lines of code. Hence if we need to determine $iter$ from the code, we do not require sophisticated techniques involving static analysis [16]. The other input variables, i.e., number of nodes n and input dataset size s, are also directly passed to the model as runtime arguments.

While modelling the estimated total completion time for the target job, the user provides an estimated upper bound for the number of iterations $iter$ for the target job, as an input to the model. During the actual running time of the target job, the user provides the number of iterations $iter_{exec}$ as a runtime argument to the job [5]. The number of iterations $iter_{exec}$ provided in the running time may differ from the number of iterations $iter$ provided in the modelling phase. The difference between $iter_{exec}$ and $iter$ may cause: 1) unpredicted wastage of cluster resources, and 2) the failure to satisfy the SLO. In that case, the estimations need to be redone, with a new input value for the number of iterations. For multiple runs of the target job with different values of the runtime argument $iter_{exec}$ supplied by the user in each run, the maximum of the $iter_{exec}$ values, i.e., $iter_{exec}^{max}$, is supplied as the new input for the estimation. The estimation using the new value $iter_{exec}^{max}$ amounts to computing the value of T_{Est} from the Equation 8 with a time complexity of $\Theta(1)$ (since the degree of T_{Est} is 1 [10]), thus incurring negligible overhead.

B. Formulation of the Model

OptEx decomposes the job completion time into four phases (Figure 1), and models the total job completion time T_{Est} as the sum of the lengths of the component phases. Thus,

$$T_{Est} = T_{Init} + T_{prep} + T_{vs} + T_{comp}, \tag{3}$$

where T_{Init} is the length of the initialization phase, T_{prep} is the length of the preparation phase, T_{vs} is the length of the variable sharing phase, and T_{comp} is the length of the computation phase. As discussed in Section III-B, the execution phase of a given Spark job comprises a permutation of low-level unit RDD tasks. The number of unit RDD tasks n_{unit} increases monotonically with increasing the input dataset size s and the number of iterations $iter$ [5]. Hence, the number of unit RDD tasks n_{unit} can be expressed as a function comprising the following terms:

- The size of the input dataset is denoted as s.
- The number of iterations in the job is given as $iter$.
- The baseline term for the number of unit RDD tasks is given as $n_{unit}^{baseline}$. It is obtained from the job profile (Section III-C). Spark enables parallel execution by dividing the input dataset into partitions, and distributing the partitions/slices among the worker nodes [5]. $n_{unit}^{baseline}$ directly corresponds to the number of partitions that the input dataset is comprised of. The number of partitions can be: 1) computed from the size s of the input dataset and the number of iterations $iter$ [5], or 2) programmatically provided as a parameter to the built-in transformation method used to create the RDDs from the input dataset [5].
  For example, the Spark Wordcount program, working on input files from a HDFS backend, divides the input dataset into as many partitions as the number of HDFS blocks comprising the input files. Consider a Wikipedia dump [17] consisting of 164 files, where the size of each file is less than the HDFS block size. Hence the number of partitions, and in turn the number of unit RDD tasks is 164. Thus, the baseline $n_{unit}^{baseline}$ is 164.

Thus, the increase of n_{unit}, with respect to the above baseline $n_{unit}^{baseline}$, in terms of the parameters s and $iter$, is expressed as

$$n_{unit} = n_{unit}^{baseline} \times s \times iter. \tag{4}$$

As discussed already in Section III-C, the length of the initialization phase T_{Init} and the length of the preparation phase T_{prep} are directly estimated from the corresponding components in the job profile (Section III-C). As discussed in Section II, the length of the variable sharing phase T_{vs} and the length of the computation phase T_{commn} vary with respect to the input variables. Hence the length of the variable sharing phase T_{vs} (Equation 1) and the length of the computation phase T_{comp} are estimated as functions of the job profile components, and the input variables (Section III-C). The expression for the length of the variable sharing phase T_{vs},

comprising the baseline value $T_{vs}^{baseline}$ and coefficient $coeff$ obtained from the job profile, is given by Equation 1.

The length of the computation phase T_{comp} in Equation 3 can be further decomposed into the following two logical components: **A)** $\boldsymbol{T_{commn}}$: The length of the communication phase T_{commn} is obtained from the Equation 2. **B)** $\boldsymbol{T_{exec}}$: The length of the execution phase T_{exec} in Equation 3 comprises the actual execution of k RDD operations comprising the job on the worker nodes (Section III-C). T_{exec} depends on various factors [5]: 1) the running times of the unit RDD tasks comprising the given job, 2) the number of iterations $iter$, 3) the number of stages in the job, 4) parallelization of the job across the worker nodes, and 5) sharing of the RDD variables across the cluster. Hence, execution phase length T_{exec} is expressed as the sum over the estimated computation times of all unit RDD tasks comprising j, along with coefficients accounting for the above factors. Thus, the length of the execution phase T_{exec} of job a, without taking into account the parallelization factor n, is given as:

$$T_{exec} = iter \times \sum_{k=1}^{n_{unit}} M_a^k, \tag{5}$$

where n_{unit} is the number of unit RDD tasks given in Equation 4, M_a^k is the average job execution time of a unit RDD task k comprising the job a, and $iter$ is the number of iterations in the job.

Following prior work on modelling execution of parallel tasks [7], the overall length of the computation phase T_{comp} is divided by the factor n, taking into account the parallelization of the across the n worker nodes. Thus, the computation phase is rewritten as the sum of its two components, divided by n:

$$T_{comp} = \frac{T_{commn} + T_{exec}}{n}. \tag{6}$$

Combining the Equations 2, 5, and 6, we get

$$T_{comp} = iter \times \sum_{k=1}^{n_{unit}} M_a^k/n + \frac{A \times s}{n}, \tag{7}$$

where n_{unit} is the number of unit RDD tasks given in Equation 4, and $A = \frac{cf_{commn} \times T_{commn}^{baseline}}{s_{baseline}}$.

TABLE III: Stepwise Accuracy of Estimations

$iter$	n	T_{vs}(sec)	T_{commn}(sec)	T_{exec}(sec)	T_{comp}(sec)	T_{Est}(sec)	T_{Rec}(sec)
5	5	1.5	18	16	34	68.52	78
5	10	3	9.88	8	17.88	53.88	72
5	15	4.5	9.5	4	13.5	51	66
5	20	6	9.3	2	11.4	50.4	54
10	5	3	28.2	24	52.2	88.2	96
10	10	6	7.74	12	19.74	58.74	72
10	15	9	5.4	6	11.4	53.4	66
10	20	12	3	3	6	51	54
15	5	4.5	37.9	32	69.9	107.4	108
15	10	9	8.3	16	24.6	63.6	78
15	15	13.5	5.7	8	13.7	60.7	72
15	20	18	2.4	4	6.4	57.4	60
20	5	6	40.2	48	88.2	127.2	114
20	10	12	12.2	24	36.2	81.4	84
20	15	18	8.5	12	17.5	68.5	72
20	20	24	6.2	6	12.2	68.52	60

Finally, combining the Equations 1 and 7 in Equation 3, the estimated total completion time for the target job is given as

$$T_{Est} = T_{Init} + T_{prep} + n \times iter \times C + iter \times B/n + \frac{A \times s}{n}, \tag{8}$$

where n_{unit} is the number of unit RDD tasks given in Equation 4, $A = \frac{cf_{commn} \times T_{commn}^{baseline}}{s_{baseline}}$, $B = \sum_{k=1}^{n_{unit}} M_a^k$, and $C = coeff \times T_{vs}^{baseline}$. Table III shows the stepwise calculations for the length of the various phases in the estimated total completion time T_{Est} for the MovieLensALS application, in standalone mode, with varying number of nodes n, and the number of iterations $iter$, on m1.large Ec2 instance.

V. Estimation of Cost Optimal Cluster Composition

OptEx models the completion time T_{Est} (Equation 8) of a Spark job on a cluster comprising virtual machine instances provisioned from a cloud service provider, like Amazon (EC2), RackSpace, Microsoft, etc. The OptEx model is further used to estimate the cost optimal cluster composition for running a given job on virtual machine instances provided by any cloud provider, under the job completion deadline specified in the SLO, while minimizing the service usage cost. Let the optimal cluster size be given as $n = \sum_{t=1}^{m} n_t$, where n_t is the number of virtual machine instances of type t (depends on the instance offerings of the chosen cloud provider), and m is the total number of possible machine instance types. Let total service usage cost of running the given job on the cloud be denoted by C. Let c_t be the hourly cost of each machine instance of type t (depends on the current rates charged by the chosen cloud provider), and T_{Est} be the estimated completion time of the given job (Equation 8). Our objective is to determine the cost optimal cluster composition for finishing the given Spark job within an SLO deadline with minimum service usage cost. This goal can be mathematically stated as: optimize the objective function

$$C = \sum_{t=1}^{m} c_t \times n_t \times T_{Est}, \tag{9}$$

and obtain the cost optimal cluster composition, given as N_t =

$$\{n_t \,|\, 1 \leq t \leq m\},$$

under the constraint $T_{Est} < SLO$, where SLO is the given deadline, and T_{Est} is estimated using Equation 8.

We optimize the above objective function (Equation 9) and determine an optimal cluster configuration given by N_t, under the constraint $T_{Est} < SLO$. The above constraint involving T_{Est} is described as a convex nonlinear function over n (Equation 8), and is twice differentiable with respect to n, i.e., both first and second derivatives of T_{Est} exist with respect to n. The above optimization problem, for minimizing the cost C, under the nonlinear constraint $T_{Est} < SLO$ is solved using the Interior Point algorithm [10]. The solution to the above optimization problem enables: 1) estimating whether a given job will finish under the deadline SLO, 2) optimal job scheduling under the given deadline SLO, while minimizing cost C, and 3) estimating optimal cluster composition, given a cost budget C and an SLO.

VI. Implementation and Evaluation

A. Experimental Setup

The experimental setup consists of Apache Spark version 1.2.1, built-in within the Cloudera Express 5.3.1 package, on

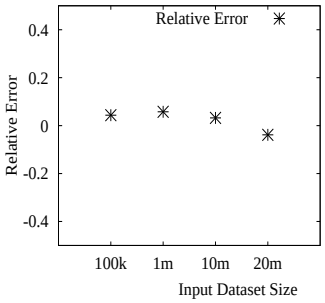

(a) Accuracy of Estimation for ALS With Varying Input DataSet Size in stand-alone mode with 20 worker nodes

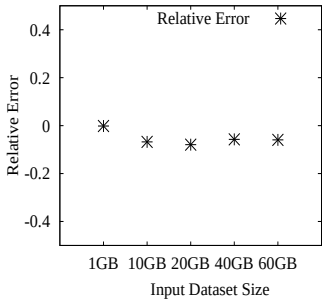

(b) Accuracy of Estimation for PageRank With Varying Input DataSet Size in stand-alone mode 20 worker nodes and 5 iterations

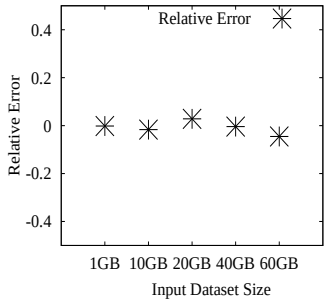

(c) Accuracy of Estimation for WordCount With Varying Input DataSet Size in stand-alone mode with 20 worker nodes

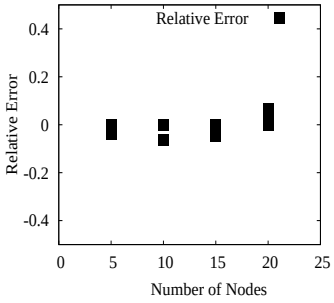

(d) Accuracy of Estimation for PageRank With Varying Number of Nodes in stand-alone mode

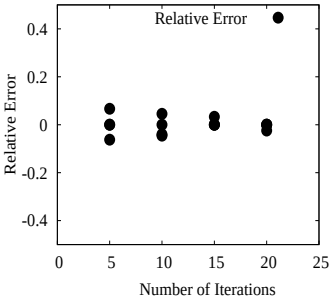

(e) Accuracy of Estimation for PageRank With Varying Number of Iterations in stand-alone mode

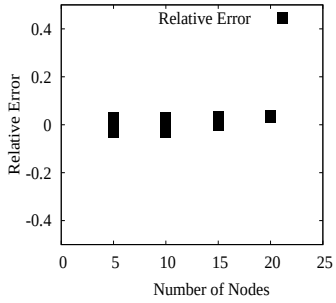

(f) Accuracy of Estimation for Logistic Regression With Varying Number of Nodes in stand-alone mode

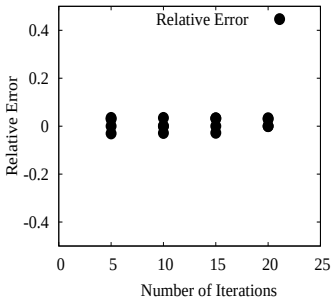

(g) Accuracy of Estimation for Logistic Regression With Varying Number of Iterations in stand-alone mode

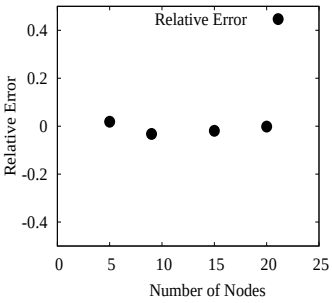

(h) Accuracy of Estimation for WordCount With Varying Number of Nodes in stand-alone mode

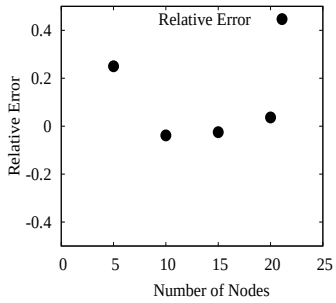

(i) Accuracy of Estimation for WordCount With Varying Number of Nodes in YARN mode

Fig. 2: Accuracy Of Estimations against varying input dataset size, number of nodes and iterations

a cluster of m1.xlarge Amazon EC2 machine instances, each comprising 8 cores, 15 GB of RAM, and 10 GB EBS, and running RedHat Enterprise Linux version 6. We use HDFS as the backend for storing and processing the input dataset. The underlying Hadoop cluster has a namenode, a secondary namenode, and 5 datanodes under a replication factor of 3.

B. Experimental Procedure

We use the Interior point algorithm [10] from the Optimization toolbox of the Matlab version 2013b for solving the given non-linear convex optimization problem (Section V), and estimating the cost optimal cluster composition. We use the default FIFO scheduler. The average scheduler delay is 4 ms, and can be neglected relative to the other components of the execution time. The input workload for the MovieLensALS application is the 10-M MovieLens dataset obtained from grouplens.org [14]. PageRank is evaluated with the social network dataset for LiveJournal [17], an online community comprising roughly 10 million members. The LiveJournal dataset has 4847571 nodes and 68993773 edges. The input workload for the Wordcount application are the Wikipedia dumps obtained from SNAP [17].

C. Technique for Generating Job Profiles

For computing the job profile of a target given job j with respect to a virtual machine instance type t, we run the representative job a (Section III-B) for the application category corresponding to the target job j (Section III-A), on a cloud

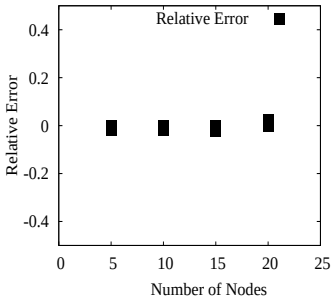

(a) Accuracy of Estimation for PageRank With Varying Number of Nodes in YARN mode

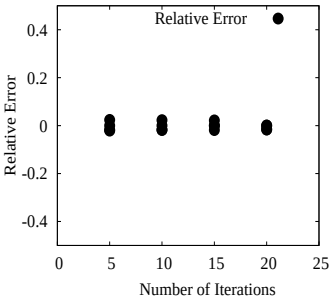

(b) Accuracy of Estimation for PageRank With Varying Number of Iterations in YARN mode

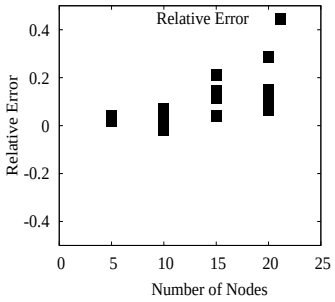

(c) Accuracy of Estimation for Logistic Regression With Varying Number of Nodes in YARN mode

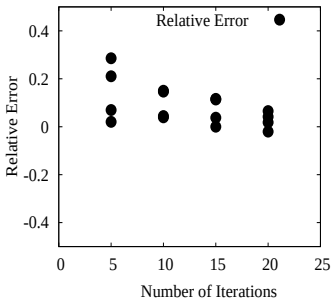

(d) Accuracy of Estimation for Logistic Regression With Varying Number of Iterations in YARN mode

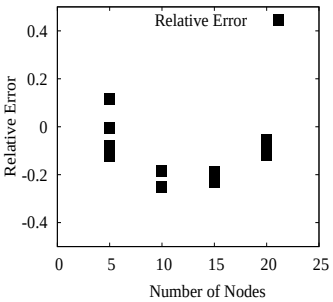

(e) Accuracy of Estimation for ALS With Varying Number of Nodes in stand-alone mode

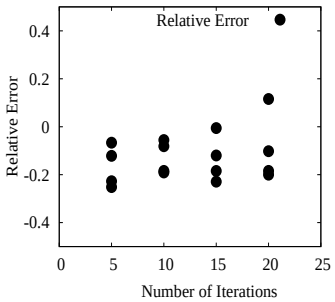

(f) Accuracy of Estimation for ALS With Varying Number of Iterations in stand-alone mode

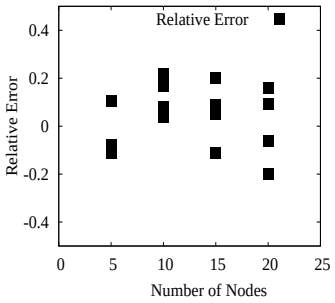

(g) Accuracy of Estimation for ALS With Varying Number of Nodes in YARN mode

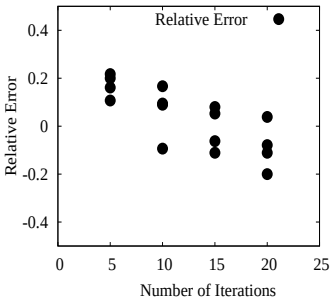

(h) Accuracy of Estimation for ALS With Varying Number of Iterations in YARN mode

Application	Stand-alone	YARN
PageRank	0.019	0.012
Wordcount	0.036	0.087
LR	0.020	0.086
MovieLensALS	0.094	0.126
Mean	0.06	
Standard Deviation	0.04	
Variance	0.0016	
Range	0.012 - 0.13	
Majority Bound	0.02	
Confidence Interval	0.056-0.062	

(i) Mean Relative Error of OptEx Estimations

Fig. 3: More Accuracy Results and the Observed Mean Relative Error Sattistics

instance of type t in stand-alone mode with a benchmark workload [18]. We estimate the components of the job profile from the snapshots of the execution flow of the representative job obtained using YourKit Java Profiler [19]. To minimize overhead, YourKit is run in the sampling mode.

D. Accuracy of the Estimations Using OptEx

Being the first work in modelling Spark jobs, OptEx has no prior baseline to compare with. However, we demonstrate (see Figures 2 and 3) that OptEx provides accurate (i.e., average relative error 0.06) estimations of the job completion time against variations in all the input parameters of the model (i.e, against increasing size of dataset, number of nodes, and the number of iterations), and on applications of different categories. From the estimated completion time T_{Est} and the recorded (i.e., observed) completion time T_{Rec}, we compute the relative error $RE = (T_{Est} - T_{Rec})/T_{Rec}$. The Figures 2(a), 2(b), and 2(c) illustrate the variations in the relative error RE, for the MovieLensALS, Wordcount, and PageRank applications, with increasing size s of the input dataset.

The figures 2(h), 2(d), 2(f), and 3(e) illustrate the variations in the relative error RE, for MovieLensALS, Wordcount, PageRank, and Logistic Regression, against varying cluster size n, in the stand-alone mode. The figures 2(e), 2(g), and 3(f) illustrate the variations in the relative error RE for the same applications, against varying number of iterations $iter$, in the stand-alone mode. Figures 2(i), 3(a), 3(c), and 3(g) illustrate the variations in the relative error RE for the same applications with varying n in the YARN mode. Figures 3(b), 3(d), and 3(h) illustrate the variations in the relative error RE for the same

TABLE IV: Optimal Scheduling With Estimated Optimal Cluster Size Under Varying SLO

SLO(sec)	Mode	App	Iterations											
			5			10			15			20		
			n	T_{Est}(sec)	T_{Rec}(sec)	n	T_{Est}(sec)	T_{Rec}(sec)	n	T_{Est}(sec)	T_{Rec}(sec)	n	T_{Est}(sec)	T_{Rec}(sec)
200	Standalone	ALS	3	103.2	107	5	88	93	6	89.67	92	7	88.84	95
240	YARN	ALS	2	240	250	3	228.44	235	4	211.5	215	4	240	237
350	Standalone	Wordcount	1	332.8	345									
800	YARN	Wordcount	1	794.8	821									
150	Standalone	ALS	3	100	97	4	99	97	4	140.5	143	6	103.22	113
200	YARN	ALS	3	174.22	215	4	181	185	5	178.56	186	5	198	195
330	Standalone	Wordcount	2	321.69	325									
790	YARN	Wordcount	2	783.69	781									
100	Standalone	ALS	4	79.5	75	6	76.11	78	6	89.67	91	7	88.84	94
160	YARN	ALS	4	150.5	157	5	159	158	6	160	155	7	159.84	158
325	Standalone	Wordcount	2	321.69	323									
785	YARN	Wordcount	2	783.69	782									
75	Standalone	ALS	5	68.52	65	7	75	67	8	71.88	72	9	73.1	71
140	YARN	ALS	5	139.52	138	7	139.92	135	9	138.07	133	9	139.1	135
320	Standalone	Wordcount	3	319.64	315									
783	YARN	Wordcount	3	781.64	754									
60	Standalone	ALS	7	58.96	56	10	58.76	57	11	61.1	55	12	62.56	58
121	YARN	ALS	22	121.05	118	31	121.02	115	39	120.97	112	42	121.1	113
319	Standalone	Wordcount	4	318.93	332									
781	YARN	Wordcount	4	780.93	772									

Category	Standalone				YARN			
	Standard Deviation	Variance	Mean	Confidence	Standard Deviation	Variance	Mean	Confidence
MLlib	0.373	0.139	1.106	0.183	0.702	0.493	2.2790	0.344
Spark Streaming	0.265	0.07	5.7	2.794	2.85	8.15	13.275	1.4
Spark SQL	1.96	3.85	32.63	0.96	3.263	10.65	48.88	0.087
GraphX	8.348	69.696	26.687	4.38	5.86	34.4	53	2.874

TABLE V: Confidence Interval of Estimation With Varying Choice of Representative Jobs for Each Job Category

applications against varying $iter$, in YARN mode. We evaluate the OptEx model with the mean relative error metric [10] $\delta = \frac{\sum_{j=1}^{k}|T_{Est}-T_{Rec}|/T_{Rec}}{k}$, where k is the total number of jobs submitted. The absolute differences between T_{Est} and T_{Rec} eliminate the signs in the error, and gives the magnitudes of the errors. The error values δ are given in the Table 3(i). The average δ score for all the cases is 0.06, i.e., 6%.

E. Analysis of the Results

The magnitude of the relative error RE for the experiments with Wordcount, PageRank, and Logistic Regression applications, representing the Streaming and GraphX categories (Section III-A), in stand-alone mode, is strictly within 0-0.06, (Figures 2(a) through 3(b)), bounded by a 95% confidence interval of 0.056-0.062 (Table 3(i)), except for one observation in Figure 2(i). The experiments with increasing dataset size yield relative error of magnitude between 0.007 to 0.05 (Figures 2(a), 2(b), 2(c)).

The estimated Spark job execution time T_{Est} comprises two components X_1 and X_2, where $X_1 = T_{Init} + T_{prep}$ and $X_2 = n \times iter \times C + iter \times B/n + \frac{A \times s}{n}$ (Equation 8). The first component X_1 is independent of variations in the values of the input variables. The second component X_2 comprises the last three phases of the Spark job execution (Section IV), each phase varying differently with respect to the input variables n, $iter$, and s (Section IV-A). Hence, X_2 accounts for the observed random variations in the relative error, with respect to variations in the input variables (Figures 2 and 3).

The execution phases in X_2 encompass the execution of the job stages, comprising unit RDD tasks, on the worker nodes (Section IV-B). The execution of the job stages on the workers is inherently non-deterministic (unpredictable) in nature, due to the dependency on various components of the Spark cluster, like the driver, the cluster manager, the workers, etc., [5]. The job stages may get unpredictably delayed, i.e., can fail and get retried by the master repeatedly, due to various factors like momentary unavailability of required resources, delays in allocation of resources by the master, communication delays among the workers, etc., [5]. The above unpredictable delays in the job stages, however small, can cause the observed values of X_2 to deviate randomly from the estimated values of X_2, while X_1 stays constant (Section IV-B). This, in turn, causes the overall observed completion time T_{Rec} of the job, to vary unpredictably with respect to the estimated job completion time T_{Est}, estimated from X_1 and X_2 (Section IV). This causes the observed random variations in the values of the relative error (i.e., $RE = T_{Est} - T_{Rec}$), though still bounded by the confidence interval of 0.056-0.062 (Figures 2 and 3).

The relative error increases slightly with increasing number of nodes (Figures 2(h), 2(d), 2(f), 2(i), 3(a), 3(c), 2(a), 3(e), and 3(g)). Worker nodes increasing in number augments the chances of unpredictable failures of the job stages due to dependency on communication between a larger number of nodes, causing unpredictable variations in the component X_2 of the overall job completion time T_{Est}. This, in turn, causes the observed job completion time T_{Rec} to deviate more unpredictably from the estimated completion time T_{Est}, estimated from X_1 and X_2. The result is greater variation in the relative error (i.e., $RE = T_{Est} - T_{Rec}$) with increasing number of nodes. Our goal is to provide correct estimations for SLO-driven user-facing applications. Few user-facing applications, that work under an SLO deadline, will require more than 50 nodes [20, 21]. OptEx can provide estimations for typical SLO-driven user-facing applications with a relative error close to 0 (Figures 2 and 3). Applications that do not meet this criteria are

batch processing applications, like bioinformatics, genomics, data analytics applications, etc., which typically do not work under a deadline [22].

For experiments run in YARN mode, the variations in the observed relative error, with respect to the variations in the input variables, are noticeably larger than the experiments run in stand-alone mode (Figures 3(a), and 3(c)). In YARN mode, the submitted jobs are additionally dependent on the YARN resource manager to allocate resources, and to execute the jobs on the worker [23]. Hence, the chances of unpredictable delays in the intermediate stages of a job are greater in YARN mode due to additional communication between the YARN resource manager and the Spark master [23]. Thus, the chances of observing randomness in the relative error is greater for applications run in the YARN mode, though the magnitude of the average error is 0.04. Further, the relative error, for YARN mode, is even closer to 0 for applications with number of iterations larger than 10, representing production level use cases [20, 21] (Figures 3(a), 3(c), and 3(g)).

OptEx cannot account for the non-deterministic delays in communicating the intermediate RDD objects among the worker nodes during the execution of an iterative Spark job on the workers [5]. The above delays result in deviations in the observed length of the job stages, comprising the component X_2, from the estimated completion time [5]. For experiments with large number of iterations, the job stages in the initial iterations cache the intermediate RDD objects locally in the worker nodes, resulting in a decrease in the time spent in communicating the RDD objects among the workers during the later iterations [5]. This results in a decrease in the deviation in the observed length of the job stages comprising X_2 from the estimated lengths of the stages. This, in turn, reduces the deviations in the overall observed completion time T_{Rec} with respect to the estimated overall completion time T_{Est}, estimated from X_1 and X_2. Indeed, with increasing number of iterations, a decreasing trend is observed in the relative error (Figures 2 and 3). So, we believe that OptEx can provide more accurate estimations with typical production level use cases, which typically involve number of iterations larger than 10 [20, 21].

F. Optimal Scheduling and Project Planning using OptEx

Again, being the first work in modeling the execution time of Spark jobs, OptEx has no prior results to compare directly with. The closest work is Elastisizer [6], which predicts optimal cluster composition for Hadoop, but does not address Spark. Moreover, Elastisizer over predicts, on an average by 20.1% and worst case 58.6% [6]. Since OptEx uses a closed-form to estimate the completion time, it does not suffer from over-prediction. Table IV demonstrates the effectiveness of the constrained optimization techniques of OptEx (Section V) in designing optimal scheduling strategies. For each application (refer to the 3rd column of Table IV) running in Standalone or YARN mode (the 2nd column), Table IV gives the cost optimal cluster composition (the optimal cluster size n is given in the 4th, 7th, 10th and 13th columns against 5, 10, 15, and 20 iterations, respectively) for executing a given job under given SLO deadlines (the 1st column), while minimizing the cost of usage of the virtual machine instances. The 5th, 8th, 11th, and 14th columns of Table IV give the completion times T_{Est} estimated using OptEx, for 5, 10, 15, and 20 iterations, respectively. The 6th, 9th, 12th, and 15th columns of Table IV give the recorded completion times T_{Rec} with the estimated cluster composition.

TABLE VI: Optimal Scheduling With Estimated Optimal Cluster Size Under Given Cost Budget

Budget($)	Mode	App	n	T_{Est} (sec)	T_{Rec} (sec)
0.3	Standalone	ALS	53	49.17	48
0.8	YARN	ALS	58	120.15	115
1	Standalone	Wordcount	27	318.02	321
1.5	YARN	Wordcount	16	780.05	775
0.2	Standalone	ALS	35	49.4	50
0.5	YARN	ALS	36	120.01	119
0.8	Standalone	Wordcount	22	318.03	321
1.2	YARN	Wordcount	13	780.09	780
0.15	Standalone	ALS	26	49.72	50
0.4	YARN	ALS	29	120.58	117
0.6	Standalone	Wordcount	16	318.06	311
1	YARN	Wordcount	11	780.12	757
0.1	Standalone	ALS	17	50.69	52
0.3	YARN	ALS	21	121.1	125
0.4	Standalone	Wordcount	11	318.12	315
0.8	YARN	Wordcount	8	780.24	780
0.08	Standalone	ALS	13	51.89	50
0.2	YARN	ALS	14	122.49	120
0.1	Standalone	Wordcount	5	318.6	310
0.5	YARN	Wordcount	4	780.94	780

Following [7], we propose a statistic $\mathcal{S}$ to measure the effectiveness of OptEx in estimating whether a given job will satisfy the SLO deadline, while minimizing the cost. $\mathcal{S}$ gives the percentage of cases which did not violate the SLO deadline, in the experiments recorded in the Table IV. $\mathcal{S}$ evaluates to approximately 98%, which proves that OptEx is, in fact, very effective for scheduling Spark jobs on the cloud, while minimizing the service usage cost.

Table VI demonstrates that OptEx can be used in project planning for optimal cluster provisioning under given budget, while optimizing job execution times. Table VI records the optimal cluster size (the 4th column) required to run a given application (the 3rd column) in Standalone or YARN mode (the 2nd column) estimated using Equation 8 under different values of the cost budget (the 1st column), while optimizing the completion times. The 5th column of Table VI gives the completion times T_{Est} estimated using OptEx, and the last column gives the recorded completion time T_{Rec} with the estimated cluster composition.

G. Confidence Under Varying Choice of Representative Jobs

The job completion times depend on the job profile generated using representative jobs. With the assumptions regarding the choice of application category (Section VI-A), Table V gives the mean, standard deviation, variance, and 95% confidence intervals for the estimated completion time T_{Est}, under varying choice of representative jobs. The function T_{Est} is a nonlinear function over the integer variable n, i.e., $T_{Est} = f(n)$ (Section V). Let μ and σ be the sample mean and variance of the job completion times for the experiments, under varying choice of representative job for each category, given in Table V. The standard deviation and variance (Table V) of the function represents the stability of the function f, under given variations in the choice of representative job.

The expectation and variance (Table V) is computed using Taylor expansions [10], and can be expressed as follows: E$[f(n)] \approx (\mu) + \frac{f''(\mu)}{2}(\sigma)^2$, and Var$[T_{Est}] \approx (f'(E[n]))^2$. The confidence interval acts as a tolerance bound that limits the estimated values of T_{Est} within an acceptable range. We say that as long as 95% of the estimated T_{Est} values remain within the interval $\mu - \alpha\sigma$ to $\mu + \alpha\sigma$ (Table V), the estimation is acceptable with 95% confidence level.

VII. Related Work

The use of job profiling, performance modeling [12], and benchmarking techniques [24] for efficient load balancing [2, 3], cost [4] and power optimization, have been attempted in quite a few cloud based systems. There has been considerable amount of work [6, 7, 25] on job scheduling and resource allocation on Hadoop. Verma et. al. [7] presents the design of ARIA, a framework for optimal resource allocation for Hadoop MapReduce. ARIA [7] models Hadoop MapReduce job execution, and cannot be readily applied to other parallel processing frameworks because: 1) it uses Hadoop-specific configuration parameters like map-reduce slots, and 2) it constructs job profile with Hadoop-specific statistics, like running time of map, reduce, and shuffle phases. Elastisizer [6] uses expensive search based or black box based techniques, that require a huge database, for provisioning Hadoop clusters, and has inherent problems like over predicting. Though Spark [5] is fast surpassing Hadoop in popularity and usage, there has not been much work in modelling Spark jobs yet.

VIII. Conclusions

OptEx models Spark job execution using analytical techniques. OptEx provides a mean relative error of 6% in estimating job completion time. OptEx yields a success rate of 98% in completing Spark jobs under a given SLO deadline with cost optimal cluster composition estimated using OptEx. OptEx can be used to estimate whether a given job will finish under a given deadline with the given resources on the cloud. It can be used to devise optimal scheduling strategy for Spark.

References

[1] A. Inc, *Amazon Elastic Compute Cloud (Amazon EC2)*. http://aws.amazon.com/ec2/#pricing: Amazon Inc., 2008. [Online]. Available: http://aws.amazon.com/ec2/#pricing

[2] S. Daniel and M. Kwon, "Prediction-based virtual instance migration for balanced workload in the cloud datacenters," 2011. [Online]. Available: http://scholarworks.rit.edu/article/985

[3] S. Imai, T. Chestna, and C. A. Varela, "Accurate resource prediction for hybrid iaas clouds using workload-tailored elastic compute units," in *Proceedings of the 2013 IEEE/ACM 6th International Conference on Utility and Cloud Computing*, ser. UCC '13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 171–178. [Online]. Available: http://dx.doi.org/10.1109/UCC.2013.40

[4] G. Singer, I. Livenson, M. Dumas, S. N. Srirama, and U. Norbisrath, "Towards a model for cloud computing cost estimation with reserved instances," *Proc. of 2nd Int. ICST Conf. on Cloud Computing, CloudComp 2010*.

[5] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica, "Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing," in *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*, ser. NSDI'12. Berkeley, CA, USA: USENIX Association, 2012, pp. 2–2. [Online]. Available: http://dl.acm.org/citation.cfm?id=2228298.2228301

[6] H. Herodotou, F. Dong, and S. Babu, "No one (cluster) size fits all: Automatic cluster sizing for data-intensive analytics," in *Proceedings of the 2Nd ACM Symposium on Cloud Computing*, ser. SOCC '11. New York, NY, USA: ACM, 2011, pp. 18:1–18:14. [Online]. Available: http://doi.acm.org/10.1145/2038916.2038934

[7] A. Verma, L. Cherkasova, and R. H. Campbell, "Aria: Automatic resource inference and allocation for mapreduce environments," in *Proceedings of the 8th ACM International Conference on Autonomic Computing*, ser. ICAC '11. New York, NY, USA: ACM, 2011, pp. 235–244. [Online]. Available: http://doi.acm.org/10.1145/1998582.1998637

[8] Y. Chi, H. J. Moon, and H. Hacigümüş, "icbs: Incremental cost-based scheduling under piecewise linear slas," *Proc. VLDB Endow.*, vol. 4, no. 9, pp. 563–574, Jun. 2011. [Online]. Available: http://dx.doi.org/10.14778/2002938.2002942

[9] N. Zhang, J. Tatemura, J. M. Patel, and H. Hacigümüş, "Towards cost-effective storage provisioning for dbmss," *Proc. VLDB Endow.*, vol. 5, no. 4, pp. 274–285, Dec. 2011. [Online]. Available: http://dx.doi.org/10.14778/2095686.2095687

[10] J. Leader, *Numerical Analysis and Scientific Computation*. Pearson Addison Wesley, 2004. [Online]. Available: http://books.google.com/books?id=y-XEGAAACAAJ

[11] A. S. Foundation, "Apache spark libraries," 2015, [Online; accessed 25-June-2015]. [Online]. Available: https://spark.apache.org/

[12] P. Brebner and A. Liu, "Performance and cost assessment of cloud services," in *Service-Oriented Computing*, ser. Lecture Notes in Computer Science, E. Maximilien, G. Rossi, S.-T. Yuan, H. Ludwig, and M. Fantinato, Eds. Springer Berlin Heidelberg, 2011, vol. 6568, pp. 39–50. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-19394-1_5

[13] R. Zafarani and H. Liu, "Social computing data repository at ASU," 2009. [Online]. Available: http://socialcomputing.asu.edu

[14] *MovieLens dataset, http://www.grouplens.org/data/*, as of 2003. [Online]. Available: http://www.grouplens.org/data/

[15] AMPLAB, "Big data benchmark by amplab of uc berkeley," 2013, [Online; accessed 25-June-2015]. [Online]. Available: https://amplab.cs.berkeley.edu/benchmark/

[16] F. Zuleger, S. Gulwani, M. Sinn, and H. Veith, "Bound analysis of imperative programs with the size-change abstraction," in *Proceedings of the 18th International Conference on Static Analysis*, ser. SAS'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 280–297. [Online]. Available: http://dl.acm.org/citation.cfm?id=2041552.2041574

[17] J. Leskovec and A. Krevl, "SNAP Datasets: Stanford large network dataset collection," http://snap.stanford.edu/data, Jun. 2014.

[18] A. Pavlo, E. Paulson, A. Rasin, D. J. Abadi, D. J. DeWitt, S. Madden, and M. Stonebraker, "A comparison of approaches to large-scale data analysis," in *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '09. New York, NY, USA: ACM, 2009, pp. 165–178. [Online]. Available: http://doi.acm.org/10.1145/1559845.1559865

[19] "The YourKit Java Profiler," http://www.yourkit.com, last 2008.

[20] A. S. Foundation, "Powered by spark," 2015, [Online; accessed 02-November-2015]. [Online]. Available: https://cwiki.apache.org/confluence/display/SPARK/Powered+By+Spark

[21] Cloudera, "Category archives: Use case," 2015, [Online; accessed 02-November-2015]. [Online]. Available: http://blog.cloudera.com/blog/category/use-case/

[22] S. Pappas, "9 super-cool uses for supercomputers," 2015, [Online; accessed 02-November-2015]. [Online]. Available: http://www.livescience.com/6392-9-super-cool-supercomputers.html

[23] A. S. Foundation. (2014) Apache hadoop nextgen mapreduce (yarn). [Online]. Available: http://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/YARN.html

[24] J. Dejun, G. Pierre, and C.-H. Chi, "Resource provisioning of web applications in heterogeneous clouds," in *Proceedings of the 2Nd USENIX Conference on Web Application Development*, ser. WebApps'11. Berkeley, CA, USA: USENIX Association, 2011, pp. 5–5. [Online]. Available: http://dl.acm.org/citation.cfm?id=2002168.2002173

[25] M. Alrokayan, A. Vahid Dastjerdi, and R. Buyya, "Sla-aware provisioning and scheduling of cloud resources for big data analytics," in *Cloud Computing in Emerging Markets (CCEM), 2014 IEEE International Conference on*. IEEE, 2014, pp. 1–8.

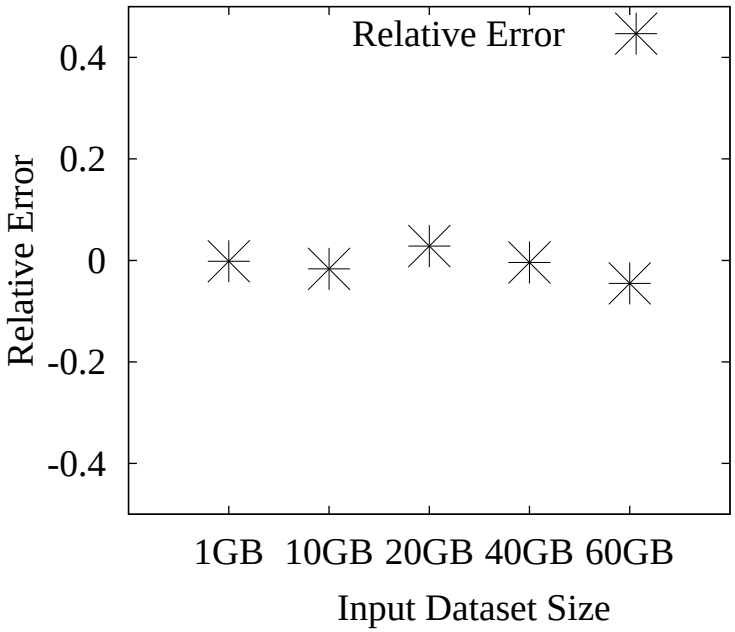
Relative Error
0.4
0.2
0
-0.2
-0.4
Relative Error
1GB
10GB
20GB
40GB
60GB
Input Dataset Size

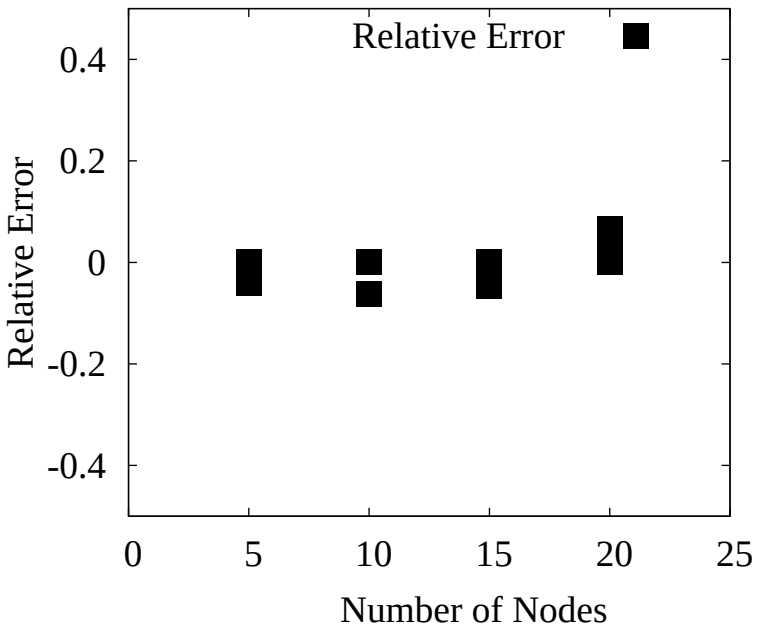
Relative Error
0.4
0.2
0
-0.2
-0.4
Relative Error
0
5
10
15
20
25
Number of Nodes

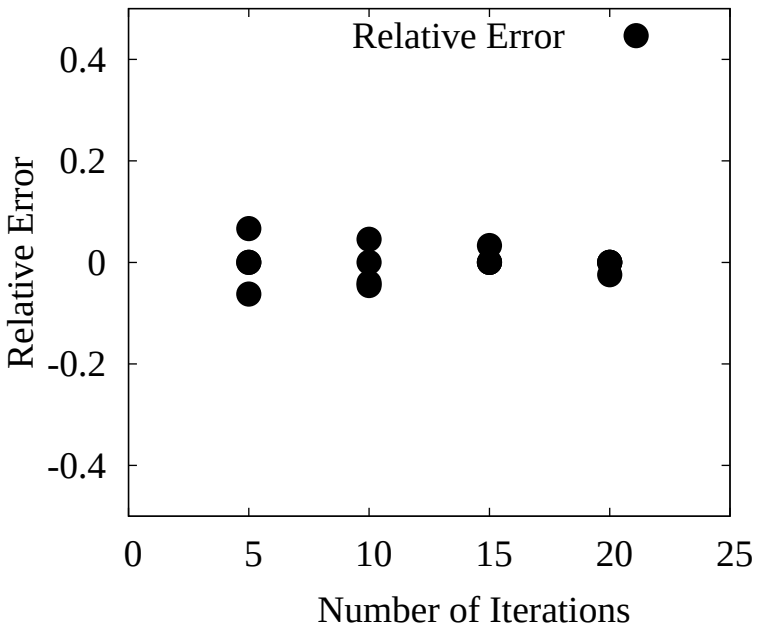
Relative Error
0.4
0.2
0
-0.2
-0.4
Relative Error
0
5
10
15
20
25
Number of Iterations

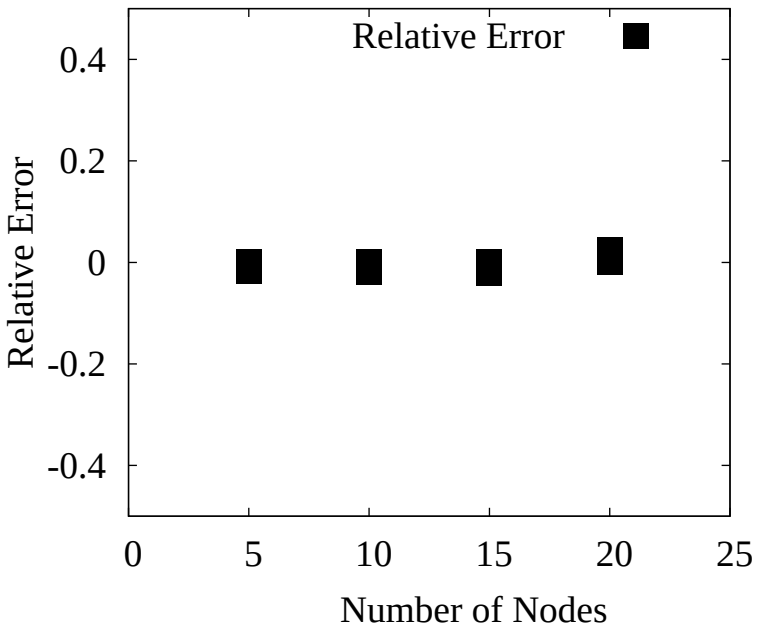
Relative Error
0.4
0.2
0
-0.2
-0.4
Relative Error
0
5
10
15
20
25
Number of Nodes

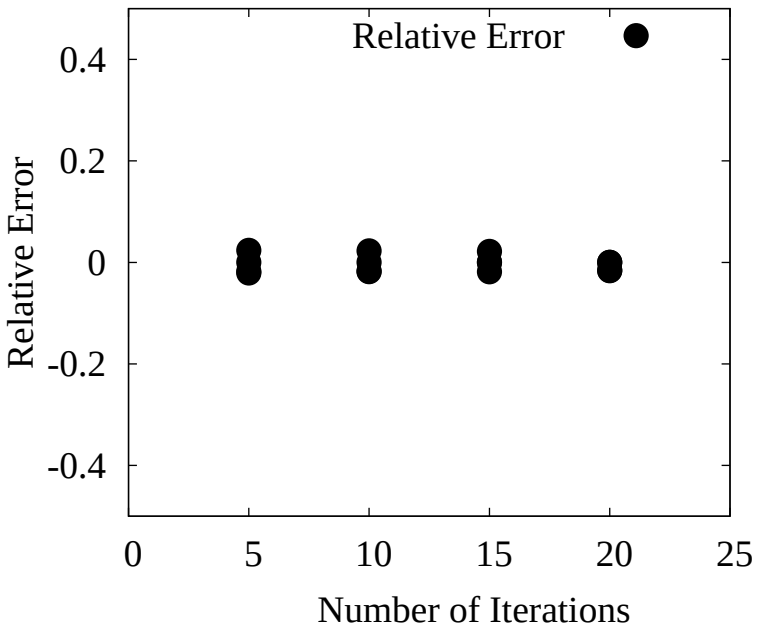
Relative Error
0.4
0.2
0
-0.2
-0.4
Relative Error
0
5
10
15
20
25
Number of Iterations